CHAPTER 5

Cryogenics

Why Cryogenics?

The niobium-titanium alloy used in the Tevatron magnet coils is only superconducting when it is at a temperature of a few degrees above absolute zero. To complicate matters further, the threshold of superconductivity drops even lower when current is present in the coil, or if it is permeated by a strong magnetic field. It would be counterproductive to try to eliminate the latter two conditions, so the only way to achieve superconductivity in the magnets is to get them cold enough. (Review Chapter 2 for more detail on the conditions required for superconductivity.) For 800 GeV operations, the coil needs to be around 5K; at 1 TeV, where current and field strength are higher, the temperature of the coil needs to be close to 4K. To establish those temperatures, the magnet coils and the surrounding stainless steel collars are immersed in a slowly flowing stream of liquid helium. Helium is the only known substance that is a liquid at the desired temperatures. Oxygen, nitrogen and most other substances have long since frozen at 5K.

Liquid nitrogen, which liquifies at 77K, is also used in the magnets, but only as an outside layer of insulation for the helium. Even then, the helium and nitrogen are separated by a layer of vacuum and reflective Mylar insulation in order to keep the helium from being overheated by the nitrogen. (Again, review Chapter 2, which includes a cross-section of a Tevatron dipole.)

Heat energy always flows from warmer to cooler. That unfortunate principle of thermodynamics makes it difficult to create an environment which is several hundred degrees colder than its immediate surroundings. (The interior of a Tevatron magnet, in degrees Fahrenheit, is 400˚ below the coldest temperature ever recorded in the Chicago area.) But there are a few standard techniques which can be used to produce cryogenic liquids, including the judicious application of *compression* and *expansion.*

Compression of a gas heats the gas. Suppose that helium is compressed with a piston. Its temperature will rise—not only because some of the energy of the moving piston may be transferred to the helium, but also because the energy originally present in the gas has been pushed into a smaller volume. A small, hot volume of gas can have the same amount of heat energy as a large, cool volume. (Remember that the only temperature at which a substance has *no* heat energy is at absolute zero. Even an ice cube contains a lot of heat compared to what it might have at a lower temperature.) The heat energy can then be easily removed from the compressed gas as it cools to ambient temperature.

Once cooled, the compressed gas can be allowed to expand; now it will drop *below* ambient temperature, because its limited supply of heat energy is spread out over a larger volume.

It is possible to cool a gas even further if it is allowed to do work as it expands, as against a piston. The work, or energy, required to move the piston is taken from the gas. Since the gas is expanding at the same time as it is losing energy, the gas is more effectively cooled because there are actually two mechanisms at work. Much of the refrigeration for the Tevatron ring uses piston-driven *expansion engines.*

Sometimes in the Tevatron, the helium does not do work as it expands and cools. An example of this process, which is less efficient than that of the expansion engines, occurs when the helium passes through a narrow aperture called a *Joule-Thompson (JT)* valve.

The last major principle of cooling to be mentioned here is that of *heat exchange.* Heat exchangers take advantage of the fact that heat flows from warm to cold. A gas or liquid, which is to be cooled, is brought into proximity (but not direct contact) with gas or liquid that is already cold. Usually the two substances flow in opposite directions.

By carefully designing the sequence of compressors, expanders, and heat exchangers, it is indeed possible to cool helium and nitrogen to the temperatures necessary for operation of the Tevatron.

There are two sources for the cryogenic liquids used by the Tevatron magnets. One is the *ring* system of refrigerator buildings and compressors. The ring system includes compressor buildings at each of the "zero" locations around the Tevatron ring, and also one at SSB. The compressors pressurize the helium coming back from the magnets and send it to the 24 refrigerator buildings distributed around the ring. The pipe transporting the helium is the *discharge header,* also known as the *3" header.* The refrigerator buildings are located on the Tevatron berm, behind the "1" through "4" service buildings. Heat exchangers just outside the buildings, and expansion engines inside, liquify the helium; each house is responsible for providing helium to upstream and downstream "strings" of magnets. The strings extend approximately halfway to the adjacent houses. Helium returning from the magnet strings is dumped into the *suction header* (also known as the *8" header*). Now at low pressure, the helium gas is pulled in by the compressors and repressurized.

The second source of cold helium, and the only source of liquid nitrogen, is the *Central Helium Liquifier (CHL).* CHL is a large building off to the side of the Tevatron ring, filled with gigantic compressors, expansion turbines and dewars, all dedicated to producing a constant supply of liquid helium and nitrogen to the *transfer line.* The transfer line originates at CHL and meets the Tevatron berm at a point between the A3 and A4 refrigerator buildings. From there it goes counterclockwise, on top of the Tevatron berm, servicing each of the 24 refrigerator buildings. The transfer line supplements the liquid helium being produced by the expansion engines, and also supplies the liquid nitrogen needed by the heat exchangers and magnets.

Since helium is constantly being fed into the Tevatron ring from CHL, sooner or later an equivalent amount of helium has to be returned to CHL. It returns in the form of excess high pressure gas from the 3" header.

All of the standard techniques described in the previous section—compression, expansion, and heat exchange—are put to use in the refrigerator buildings. The normal life cycle of the helium in the ring is summarized in the steps below. More details will follow later in the chapter. The numbers in the sequence correspond with those in Fig. 5.1:

1) Compressors around the ring take up helium from the suction header, pressurize it, and fill the discharge header. The helium, which has been heated by the compression, is heat exchanged with water in order to bring the temperature down to approximately room temperature.

2) Pressurized helium from the discharge header is again cooled, this time in a heat exchanger located outside the refrigerator building. The cold counterflow may be liquid nitrogen from CHL, cold helium returning from the magnet strings, or even a portion of the pressurized helium that has been cooled in an expansion engine.

3) The cold pressurized helium is allowed to expand, cool further, and liquify. This is usually done through an expansion engine (the *wet engine*), but a JT valve is available as a backup.

4) The liquid is cooled further by heat exchanging with liquid helium in a *dewar.* The dewar itself is filled with liquid helium from CHL and cold helium returning from the magnet strings.

5) Just before it enters the magnet string, in the *feed can*, the liquid helium is heat exchanged with the cold helium returning from the magnet string.

6) Splitting between the upstream and downstream strings, the liquid helium makes its first pass through the magnet strings, immersing the coil and the stainless steel collars. Here the helium is referred to as *one-phase* or *single-phase*, the "phase" in this case being liquid.

7) At the end of each string, the liquid helium passes through the *magnet JT* valve, where it expands, thus cooling further, and is partially converted to a gas. The magnet JT valves are located at the *turnaround boxes* at the ends of each string.

8) The expanded helium, now known as *two-phase* (liquid plus gas), is used as a counterflow to cool the one-phase further. The two-phase circuit is a shell just outside the one-phase. Flow is back toward the feed can.

9) Upon reaching the feed can, the two-phase helium is nearly boiled off and can now be considered a cold gas. It is used to heat exchange with the incoming gas in the feed can, cool the dewar, and finally cool the pressurized helium in the heat exchanger.

10) Beam energies greater than 900 GeV require colder temperatures in order to maintain the magnets in a superconducting state. A *cold compressor,* located downstream of the dewar, pumps down on the two-phase helium to lower its pressure, and thus its temperature. The colder two-phase lowers the temperature of the dewar and the magnets appropriately. The discharge of the cold compressor passes through the heat exchanger; although it is heated some by the compression, it is still cold enough to be effective as counterflow to the incoming helium.

11) After extracting as much heat as possible from the incoming gas in the heat exchanger, the low pressure helium enters the suction header, finds its way to the compressors, and is repressurized. The cycle starts over.

12) Throughout this process, CHL provides liquid nitrogen for use in the heat exchanger and the magnets, in addition to its role mentioned earlier of providing liquid helium.

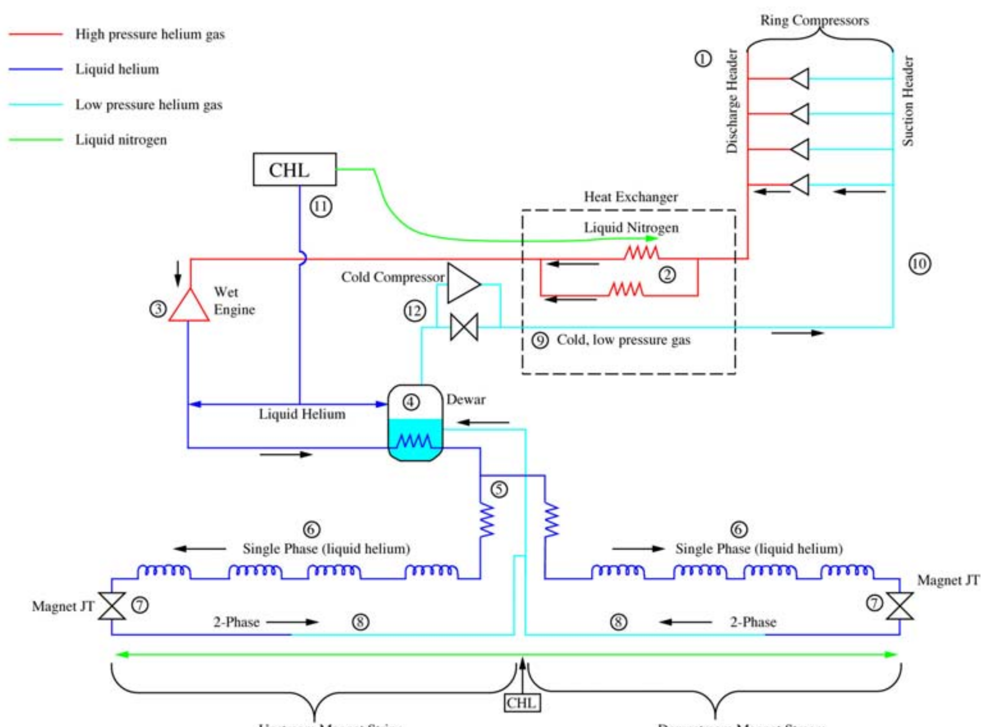

Fig. 5.1 Cryogenic Overview
See text for details

Instrumentation

Before describing the components of this scheme in more detail, it will be useful to look at some of the methods used to measure pressure and temperature.

Pressure is usually measured in *psig*. The psig unit needs to be distinguished from the slightly more familiar unit *psi* (pounds per square inch). The *g* stands for "gauge," i.e. the pressure transducer itself. Since pressure gauges are often calibrated to read zero pounds at atmospheric pressure, a measurement in psig represents the value *above* the standard atmospheric pressure of approximately 15 psi. This distinction is important in cryogenics, because if the psig of a helium circuit goes negative, it is "sub-atmospheric;" it is possible that freezable gases such as oxygen and nitrogen could leak in and clog up the plumbing. Parameter names for pressure gauge readbacks usually include "PI" (pressure indicator) as part of the name. Occasionally units of pressure are given in psi, or *psia* (the *a* standing for "absolute.")

There are three ways of measuring temperatures in the cryogenic systems. One is known as a *VPT (vapor pressure transducer*). *Vapor pressure* is actually measured in the vicinity of a liquid. In every liquid, there are at least a few molecules which have enough energy to break free of the liquid; these molecules are in a gaseous state and are responsible for the vapor pressure. The higher the temperature, the greater the number of molecules that are in the vapor. If the temperature reaches the boiling point of the liquid, the vapor pressure equals the pressure confining the liquid, and all of the liquid begins to convert into a gas.

The VPT's are actually liquid-filled bulbs fitted with pressure transducers. Although they read back in units of psig, the vapor pressure actually represents a temperature. The bulbs may be filled with nitrogen, hydrogen, neon, or helium—the choice depends on what temperature range is desired (of course, it would not work for the liquid to freeze solid in the environment it is supposed to be monitoring). Tables relating temperature to vapor pressure can be found in the parameter page HELP files. VPT parameter names usually include the phrase "TI" (for Temperature Indicator).

The second type of "thermometer" is the *carbon resistor.* These carbon resistors are simply precision versions of those used on any electronics board. Since the resistance varies with temperature, a constant current source upstairs is fed into the resistor; the voltage across the resistor can be converted into resistance, and the resistance can be scaled to temperature. (The scaling is not as obvious as one might think. In the cryogenic temperature range, the resistance of a carbon resistor is *inversely* proportional to temperature, and nonlinear at that. Fortunately, resistance values are normally converted to temperature values in the database before they are read on a parameter page.) Carbon resistors are at their best at the coldest temperatures (<10K), but very inaccurate at room temperature. Carbon resistors include "TR" (Temperature Resistor) in their name.

Knowing where each carbon resistor is located is essential when cooling a magnet string (Fig. 5.2 (a), (b), and (c)). To a first approximation, the resistors are named after the nearest quadrupole—using the E2 strings as an example, the resistor at the base of the feed can is in the spool piece next to the quadrupole at E25; it is named TRQ5. Going upstream, the numbers count back to TRQ1, in the turnaround box; downstream, TRQ9 is in the last spool piece in the string, at E29. In the case of the “1” houses, the numbers count up to TRQ0. (The downstream string at E1 continues to the spool piece at E21, so “0” really means “10.”)

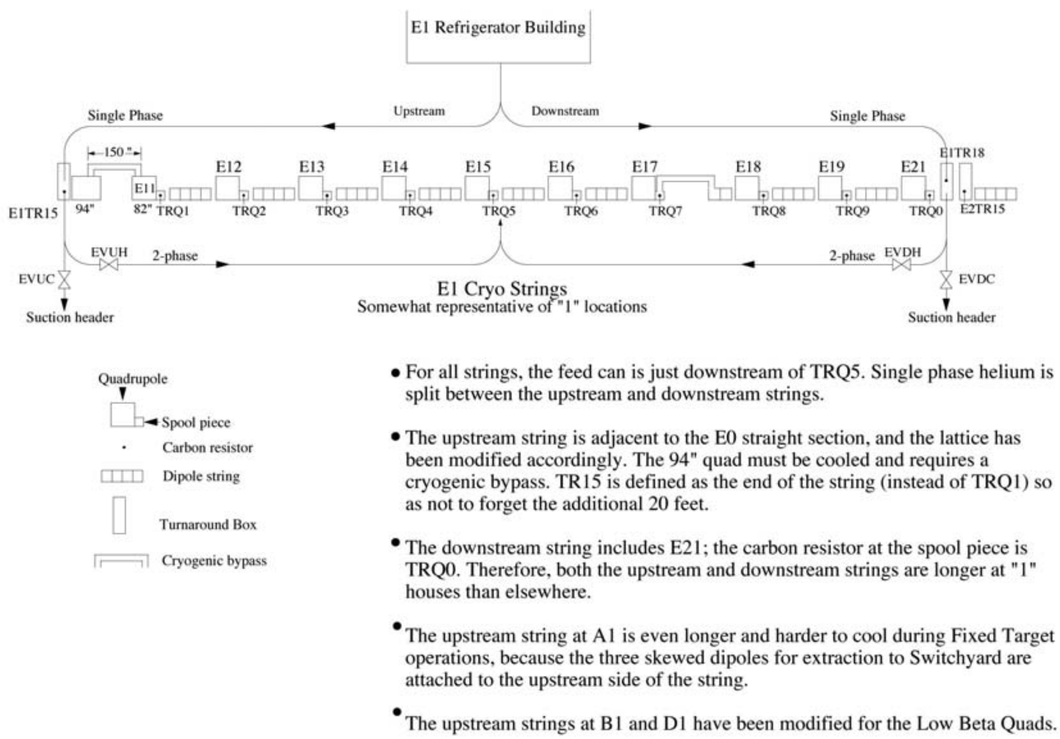

Fig. 5.2 (a) Cryogenic Layout of a "1" House

The turnaround boxes themselves are instrumented with TR15 (upstream) and TR18 (downstream). Sometimes TR15 or TR18 should be used to evaluate temperature at the end of the string. In the “1” houses, TRQ1 is separated from the turnaround box by two quads and a cryogenic bypass—or at B1 and D1, by low beta quads—so TR15 should be used instead. At the “4” houses, TRQ9 is separated from the turnaround box by a 94” quad. Another important carbon resistor is TR12, which measures the temperature of the helium entering the heat exchanger.

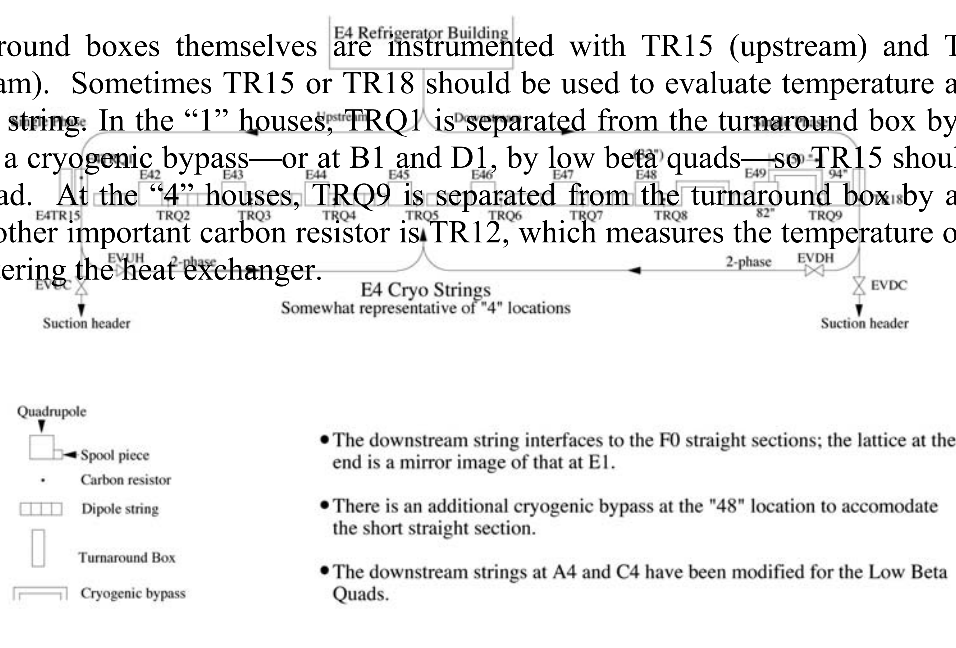

Fig. 5.2 (c) Cryogenic Layout of a "4" House

For measurements in the warmer regions of magnet strings, *platinum* resistors are used. They operate on the same principles as the carbon resistors, except that they are at their most accurate near room temperature. "TP" is a part of every platinum resistor's name.

Control Loops

There are a few generalized comments about control loops which should be made before moving on to the specific refrigerator components. The refrigerator system is a large, complex, and dynamic one, and the machinery must constantly be adapting to changing conditions. Most of the valves and engines are operated by feedback loops; they adjust their output, within minimum and maximum bounds, to try to match a measured value to a set point. A change in one component of the system may prompt a response from another component. The following discussion will emphasize the way that these variables interact in a steady-state operational mode, but necessary complications will be have to be introduced later.

There are *Finite State Machines (FSM's)* which can take control of the loops. The FSM's are also used to control other devices, as explained later.

Ring Compressors

The ring compressors pull in low pressure helium exhaled by the magnet strings, and pressurize it so it can be cooled by the expansion engines and JT valves in the refrigerator buildings. There are compressor rooms at all of the zero locations (A0, etc.) and one at SSB. Each compressor building normally houses four compressors (except for SSB, where there are only two.) If there are more than four, the additional compressors are grouped separately, at least from a naming and controls standpoint. For example, at B0 there are 8 compressors; four belong to the "B0" group and the other four belong to the "BA" group. Altogether there are 32 compressors; however, in Collider Mode one compressor at B0 is dedicated to the cryogenic solenoid at CDF and is unavailable as a ring compressor. All of the helium compressors are in parallel; that is, they all draw low pressure helium from the same set of suction headers, and, after pressurizing it, send it to a common 3" header.

By the time the helium has arrived in the suction header, it has completed its circuit of the expansion engines and magnet strings, and dropped to a pressure of about 1.2 psig.

The 1.2 psig helium from the suction header is pulled in to the compressors and leaves at a discharge pressure in the neighborhood of 280 psig. The high pressure gas enters the 3" header, which runs parallel to the transfer line on the Tevatron berm. The expansion engines in the refrigerator buildings use the compressed helium. It is important to note that even the ideal suction and discharge pressures are expected to vary somewhat from point to point around the ring. For example, there is a *suction profile* which requires different sections of the suction header to be at slightly different pressures.

There are valves in the B0 Compressor Room for shunting excess helium in the discharge line back to CHL, more about those later.

A few parameter names should be mentioned at this point. Suction pressure, which is measured at each zero building, takes the name T:xxPI1, where *xx* represents the location (i.e. B0PI1). Discharge pressure takes the name T:xxPI2.

Digital status for the individual compressors is organized under the parameter type T:xxHPyy, where *yy* is a number related to the compressor number at a given house. For some totally inexplicable reason, Compressor #1 is given the name T:xxHP30, #2 is T:xxHP40, #3 is T:xxHP50, and #4 is T:xxHP60. This strange scheme is also reflected in some of the other compressor parameter names that will be introduced later.

The ring compressors are of a *screw-type* design, not piston-driven as one might expect. Each stage consists of two large rapidly rotating screws which are roughy parallel to each other. Bridging the two screws is a *slider* (no relation to White Castle), which acts something like a zipper. If the slider is at one end of the stage, the screws are "unzipped;" they are misaligned enough that the helium gas is not forced through the system. If the slider is at the other end, the two screws mesh and the helium, trapped inside the spaces, is compressed as it passes through. The slider can also be at any intermediate position, is which case only some of the helium is compressed. The degree to which the slider causes the screws to mesh is called *loading,* and a compressor which is processing the maximum possible quantity of helium is said to be *fully loaded.*

At each individual compressor, the helium is compressed in two stages in order to keep pressure differentials at a minimum. The first (or "low") stage takes helium from the suction header; the helium typically starts between 1 and 2 psig and is compressed to about 20 psig. The second (or "high") stage takes the 20 psig helium and compresses it to the discharge header value of 280 to 300 psig.

The parameter name of the low stage slider for Compressor #1 at a given location is T:xxHV1L; the high stage slider is called T:xxHV1H. Compressor #2 parameters would be T:xxHV2L and T:xxHV2H, etc.

Regulation loops (Fig. 5.3) control the amount of loading for each stage. The low stage regulates from PI2—the discharge pressure—and in doing so determines how much helium is pulled from suction. For example, if the loop senses that the discharge pressure is too low, it sees to it that more helium is taken from suction; a greater quantity ends up in the discharge line and the pressure in the line increases. The set point for the loop is usually in the high 200's (psig), but varies a little from compressor to compressor in order to keep the discharge profile healthy and stable.

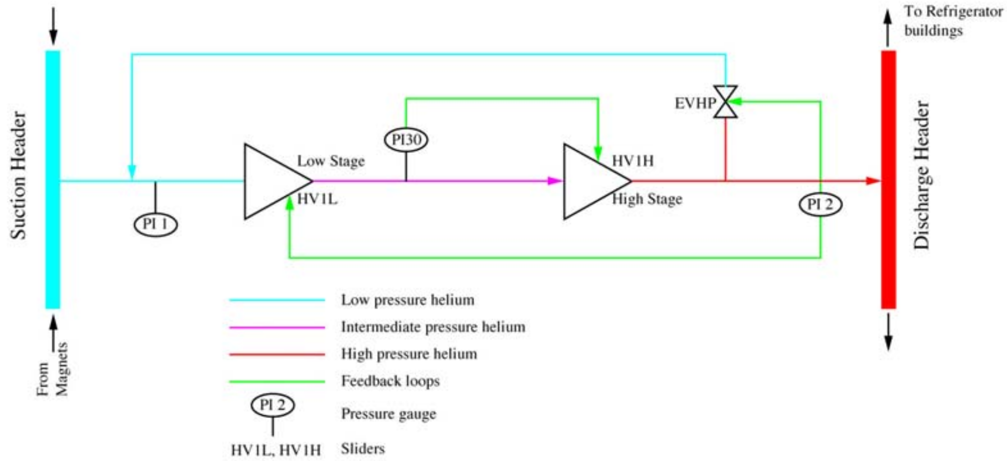

Fig 5.3 Compressor Loops

The high stage regulates from the *interstage pressure.* (The interstage pressure reading for Compressor #1 at a given location takes the name T:xxPI30.) A common set point for the interstage pressure is 20 psig. Remember that the low stage has already made up its mind as to how much helium is going to be compressed. The high stage takes that gas and compresses it. The harder the high stage works, the lower the interstage pressure. If the interstage pressure drops below the set point, the high stage senses that it is working too hard, and the slider "unmeshes" the screws a little. If the interstage pressure is too high, the slider moves to mesh the screws more tightly so that more helium is removed from the interstage area and compressed.

Compressors are somewhat messy machines—the screws have to be heavily lubricated with oil, and there is no way to prevent the oil from mixing with the helium. After compression, but before actually being allowed into the discharge line, the helium passes through separators and charcoal filters in order to remove the oil. Helium leaving the compressor is also quite hot, so it is also heat exchanged with chilled water.

T:xxEVHP, the *compressor bypass valve,* is common to all four compressors at a given location. It opens when the discharge pressure exceeds the set point; obviously, the set point is normally made higher than the output requested from any of the individual compressors. The excess high pressure gas is shunted back to the suction header.

Finally, there are two specialized valves at B0 and BA called the *kickback* valves. Remember that CHL is constantly sending liquid helium to the magnets via the transfer line. As the helium works its way through the refrigerator buildings and the magnets, it warms and vaporizes, and enters the suction lines as gas. To keep the Tevatron from exploding (just kidding, of course), kickback valves return an equivalent amount of helium to CHL from the discharge line; the valves are regulated by the suction pressure. These two valves will be examined in their full context later in the chapter when helium *inventory* is discussed.

Computer control of the compressors is an integral part of the refrigerator controls system, and will be addressed as part of that section.

From each compressor building, the discharge line connects to the 3" header located on top of the Tevatron berm. The header passes through each of the 24 refrigerator buildings, where the helium is withdrawn by the local cryogenic systems.

Because high pressure helium can present a hazard to personnel—as a mechanical energy hazard as well as from an ODH standpoint—the first valve encountered in the building is SV101, the *Emergency Shutoff Valve.* The valve and its motor control hang from the ceiling. Pushing the *crash button* on the outside wall of the building closes the valve. This action shuts down much of the equipment in the building and should only be taken judiciously.

After getting past SV101, the next stop is the heat exchanger.

Heat Exchangers

The purpose of the heat exchangers is to cool the high pressure gas as much as possible before it is sent to the wet engine. The exchangers are the long tube-shaped structures attached to the sides of the refrigerator buildings. Internally, each heat exchanger actually consists of four interconnected exchangers.

As it enters the exchanger from inside the building, the high pressure line temporarily divides into two branches (Fig. 5.4). One branch is cooled by heat exchange with liquid nitrogen supplied by CHL, and the other branch is cooled by heat exchange with the cold low-pressure helium (alias "suction") returning from the magnets in the tunnel. The loops for both exchangers look at TI5, a nitrogen VPT. TI5 measures the temperature of the low-pressure helium at the point where the two branches combine again. The strategy: If the returning helium is cold enough, use it. If it isn't, use the liquid nitrogen. If it is somewhere in between, use both.

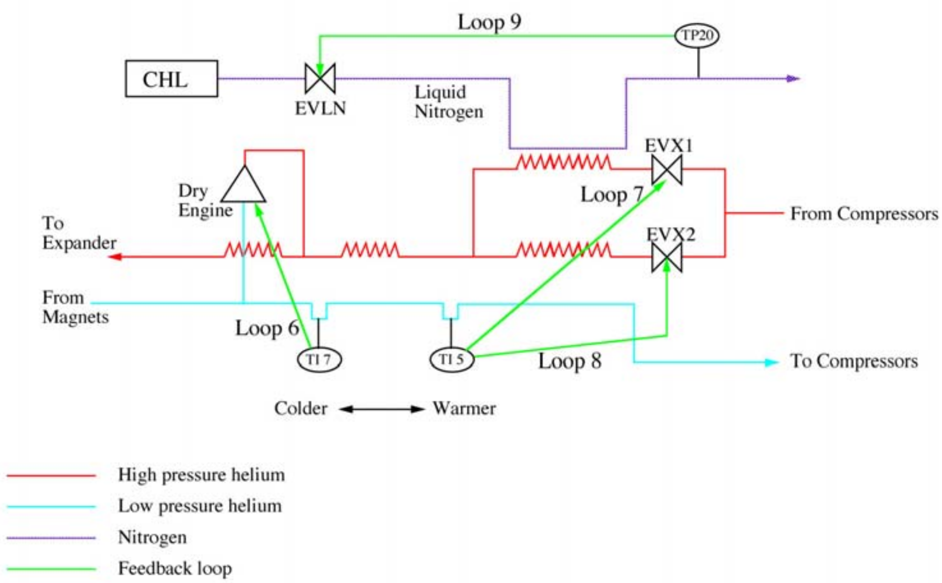

Fig. 5.4 Heat Exchanger Loops

The nitrogen exchanger is controlled by two loops—one that regulates the flow of nitrogen into the exchanger (Loop 9, EVLN), and one that controls the amount of high-pressure helium going through the nitrogen exchanger (Loop 7, EVX1). EVX1 adjusts the flow of helium through the exchanger to move TI5 toward its set point of 25 psig, which corresponds to a temperature of about 87K. Remember that the temperature of liquid nitrogen is about 77K—the valve opens up if TI5 is too warm, and closes down if it is too cold.

EVLN regulates to a platinum resistor called TP20, located in the nitrogen line downstream (nitrogen direction) of the exchanger. TP20 is assigned a set point near 200K. The purpose of EVLN's loop is to conserve liquid nitrogen; if the nitrogen is colder than 200K after heat exchanging with the helium, some of it is probably being wasted and the valve closes some.

The flow of high pressure helium through the branch that exchanges with the cold helium from the magnets is controlled by the valve EVX2 (Loop 8). Its loop is also regulated around TI5, but the feedback is opposite that of EVX1—if TI5 is too warm, EVX2 closes down to keep from warming the incoming helium. If TI5 is too cold, EVX2 opens in order to fully exploit the opportunity.

In practice, the loops are changed only during cooldown. During steady-state operation, returning helium is usually cold enough to carry the full load. EVX1 is locked nearly closed, and EVX2 is locked completely open. There is no need to conserve the low temperature of the returning helium, because this is the last point at which it will be used for heat exchange.

The heat exchanger is also equipped with an expansion engine called the *dry engine.* The dry engine, except for a smaller piston, is virtually indistinguishable mechanically from the wet engine. The only real difference between them is their location in the cooling sequence. The dry engine expands gas which is still too warm to liquify.

The dry engine is seldom used any more, but when it is, it pulls high pressure helium out of the line, expands it, and dumps it back into the suction line. The relatively cold gas then joins the helium from the magnets and helps in the heat exchange. This cooling technique is normally used when cold helium from CHL is not available.

The loop (Loop 6) for the dry engine regulates to TI7, a hydrogen VPT. Notice in Fig. 5.4 that TI7 is downstream (cold helium direction) from the dry engine. The set point for TI7 is 10 psig, which corresponds to a temperature of about 22K. The harder the dry engine works, the colder TI7 becomes. More detail about expansion engines is forthcoming in the next section.

The dry engines are not normally used because cold helium is available from CHL. The loops regulating the CHL helium are best deferred until later, when the refrigeration system as a whole has been more fully explored.

Expansion

After the heat exchanger has cooled the high pressure gas as much as possible, it is time to call on the next technique—expansion (see Fig. 5.5). In a healthy, steady-state operation, the temperature of the gas just before it enters the expanders is about 6K. Expansion will reduce the temperature to a point at which the helium begins to condense into a liquid.

There are two devices, in parallel, which can implement the expansion. One is the *wet engine.* As an *engine,* the cold gas pushes against a piston as it expands. It is doing work, exactly as steam once did in locomotive engines. And, like the steam, energy is transferred from the gas to the piston, and the gas condenses. In fact, the activity of an expansion engine is monitored by attaching a small electric generator and reading its output. The helium gas is generating power, but only enough to rob it of its energy.

The engine expels a certain quantity of helium with every stroke, so the *speed* of the engine, in RPM, is the measure of its output. The wet engine normally runs between 300 and 1200 RPM.

The feedback loop (Loop 2) regulating the speed of the engine uses a pressure gauge downstream of the output; the faster the engine runs, the greater the quantity of helium pumped, and the higher the pressure of the output. There are several pressure gauges downstream of the engine which can be used, but the preferred one is PI13. If PI13 is malfunctioning, PI14 or PI17 can be used. Since the latter two gauges are downstream of the feed can, the set point should account for the pound or so of pressure drop the helium experiences going through the feed can. The set point of PI13 during steady-state operation is 17 psig.

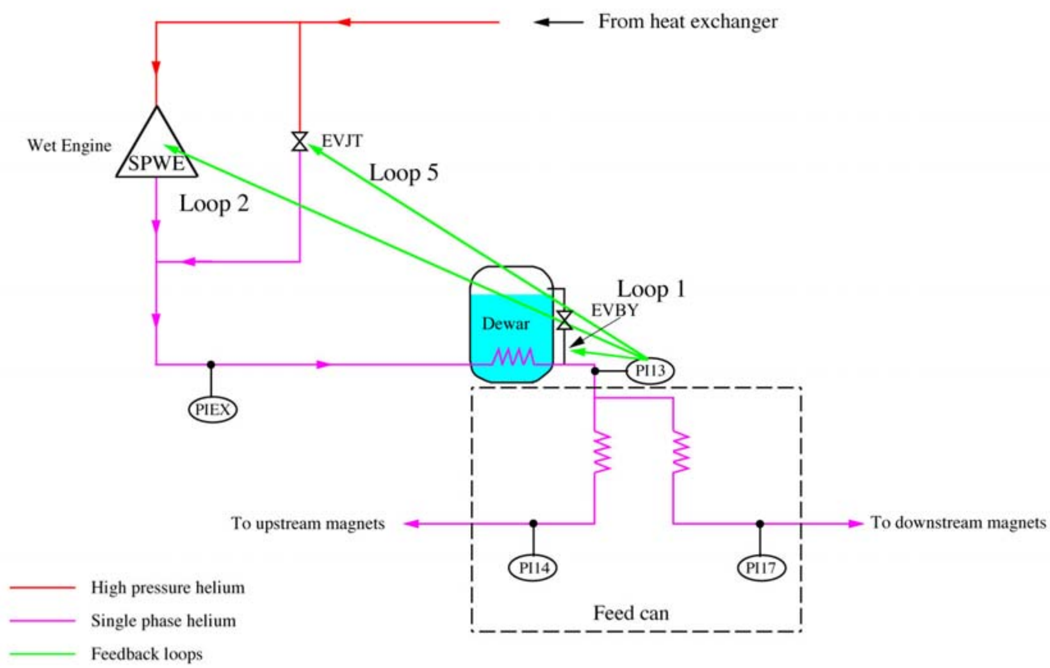

Fig. 5.5 Pressure Regulation Loops

The other expansion device is a *Joule-Thompson (JT)* valve, called EVJT. A JT valve forces the gas through a narrow aperture, after which it is allowed to expand. There is no mechanism for removing energy from the gas, so the cooling is less effective.

The wet engine is always the preferred method for liquifying the gas; the JT valve is strictly used as a backup. Not only is the quality of cooling inferior with the JT valve, but it requires a greater quantity of compressed gas to achieve the same effect. The valve, like the wet engine, regulates to PI13 (Loop 5), but the set point is 12 or 15 psig; in the event of an unexpected failure of the wet engine, the JT valve acts as a safety net. The set point of the JT can be raised if wet engine recovery is not imminent.

There is one more pressure-related valve—EVBY, the *bypass* valve—which is used during cooldowns but not usually during steady-state operations. It also looks at PI13 (Loop 1), but is set to *open* if the pressure exceeds the set point of 18 psig. It acts as a safety valve to prevent overpressurizing the magnet strings, but normally the wet engine will slow down before the pressure gets that high. During cooldowns, EVBY is used to siphon off cold helium to the dewar and heat exchanger. More about that later.

There are two additional opportunities to further cool the helium before it reaches the magnets. One is the *dewar,* which during steady-state operation is full of liquid helium. The other is in the *feed can,* where the incoming helium is heat exchanged with the 2-phase helium returning from the magnets. The 2-phase is actually colder than the single phase.

Magnet Strings

The single-phase helium from the feed can splits to supply the upstream and downstream magnet strings. Inside each magnet is the *cryostat,* which is a long tube containing the beam pipe, the magnet coil, the helium and nitrogen "circuits," and the insulating vacuum. The cryostat is surrounded by the iron yoke. In most of the Tevatron magnets, the iron yoke is at room temperature; however, the low-beta quads are of a "cold-iron" design.

The single-phase helium enters the magnets and surrounds the coil and the stainless steel collar surrounding the coil. (The entire cryogenic system was built for this one moment.) The single-phase helium slowly moves through the magnet string until it reaches the *turnaround box* at the end. During normal operations, the helium passes through the *magnet JT* valve at the turnaround box. The liquid helium, after passing through the narrow aperture, expands and cools. It is then sent back into the magnets via the 2-phase circuit, which is concentrically to the outside of the single phase circuit. Because of the JT effect, it is a few tenths of a degree colder than the single phase helium; heat exchange takes place throughout the magnet string as it returns to the feed can.

The 2-phase helium consists of liquid and gaseous helium (hence the name). Immediately after passing through the JT valve, it has a relatively high ratio of liquid to gas. As it moves back along the 2-phase circuit, it absorbs heat from the liquid helium on the inside and from the environment on the outside. As with boiling water, it remains at the same temperature even as it is being converted to gas. It is preferred that all of the liquid—but no more—boils off by the time that the end of the string is reached. In that way the temperature remains cold, but no liquid helium is wasted.

The feedback loop controlling the JT valve uses a hybrid device called a *DT* (for *differential temperature)*. The DT measures the difference between the pressure at the end of the 2-phase circuit and the temperature as measured by a VPT (DT = VPT - PI). Pressure is constant throughout the 2-phase circuit, so as long as there is a liquid/gas mixture the VPT should remain constant. If all of the liquid boils off, the temperature will begin to rise as it absorbs energy from its surroundings; in that case, the DT rises. It is even possible to "freeze out" a DT by producing a liquid below the boiling point.

The PI for the DT measurement is usually PI11, upstairs in the refrigerator building. There is a slight pressure drop going from the 2-phase circuit to PI11, so the DT usually has a set value of about 0.2 psid (*d* stands for *differential*).

Establishing the operating point for a JT valve is a fine and delicate art. Often, the value is between 50% and 60%, but it varies from string to string. The effectiveness of the JT depends on the balance between the amount of flow and the amount the valve is constricted, and that balance depends in turn on the temperature. Specifics will be discussed in the sections on cooling the strings.

The feedback loops (Loop 3 upstream, Loop 4 downstream) works by opening up the JT (increasing flow) if the DT rises above its set point, and decreasing flow by closing the JT when the DT shows that the helium is too cold. The limits are usually set to restrict changes to a few percent.

Nitrogen

Surrounding the helium layers is a vacuum break. Outside the vacuum layer is the liquid nitrogen circuit. The nitrogen is a free gift from CHL. At roughly 77K, it is an insulating layer between the helium and the balmy climate outside the magnet.

The nitrogen makes a single pass through the magnets and empties into a common *nitrogen header* behind the magnets. A pipe between A3 and A4 connects the header to CHL, where the nitrogen is reliquified.

The valves controlling the amount of nitrogen flowing through the magnets are EVUN (Loop 11, upstream string) and EVDN (Loop 12, downstream). They regulate to the temperatures measured by two platinum resistors, TP23 (upstream) and TP24 (downstream). The set point for each is 90K, slightly above liquid nitrogen temperature. The goal is to have converted all of the liquid nitrogen to a gas by the time it leaves the string, while still keeping the magnets cold.

Cryostat Vacuum

The single phase and 2-phase circuits are adjacent to each other and within a degree or so of the same temperature, but the nitrogen circuit is much warmer. The two layers of cryogens are separated by a "layer" of vacuum. The same vacuum extends to the outside of the nitrogen shields. This vacuum is completely distinct from the vacuum of the beam tube. It is also of a lower quality; beam tube vacuum is in the range of 10^{-10} torr, while the cryostat vacuum is closer to 10^{-7} or 10^{-8} torr. But it is sufficient to maintain temperatures in a cryogenically healthy magnet.

Within the vacuum layer are several layers of reflective Mylar insulation, which prevent infrared light from penetrating to the cryogenic levels. These elaborate layers of insulation work—the surface of the magnets is at room temperature, leaving no clue as to the near absolute-zero temperatures a few inches away.

Chapter 6 includes more detail on the cryostat vacuum.

Kautzky Valves
A quench which vaporizes a great deal of helium carries the risk of overpressurizing the magnets. Kautzky valves, in one of their roles, act as pressure relief valves.

Kautzkys are of a very simple design—a plastic "poppet" is held in place by a *control pressure* of 30 psig, supplied from helium bottles upstairs in the service building. (Helium, as opposed to air, is used in order to ward off contamination in the magnet string.) One side of the poppet faces the suction header, and the other side is facing one of the cryogenic circuits—single phase, 2-phase, or nitrogen. If, during a quench, the pressure inside the magnets exceeds 30 psig, the poppet opens and allows the helium to enter the suction line.

There is one Kautzky valve for each of the dipole magnets dedicated to the single-phase circuit, since that is the circuit with the greatest potential for expansion of gas. At the spool piece, there is one Kautzky valve dedicated to each of the three circuits.

The single-phase Kautzkys at the spool pieces can be remotely controlled, and during cooldowns they are sometimes opened and closed in a systematic fashion in order to expel warm helium from the magnet strings—more about that later.

The Dewar

The dewars in the refrigerator buildings were installed as part of the upgrade to high energy (1 TeV) operations. They are filled with liquid helium; during normal operations, the helium is used not only to heat exchange with the exhaust of the wet engines, but also to provide a reservoir of cold helium for the shell side of the heat exchangers. But the real reason for the existence of the dewar is to compensate for the heat load generated by the *cold compressor*, described below, during low-temperature operations.

The dewar obtains helium from several different sources, including CHL, 2-phase returning from the magnets, and, during cooldowns, single phase helium from the wet engine exhaust via EVBY.

CHL

Cold helium arrives at the refrigerator building from the transfer line; the gatekeeper at the building is EVSH, a solenoid-operated valve. It is either open or closed—almost always open when the refrigerator is operational, and closed only during maintenance periods or when CHL is not producing liquid. There are no feedback loops controlling EVSH.

Downstream of EVSH are two regulated valves controlling the flow of CHL helium: EVLH and EVQH.

EVLH (Loop 10) adds liquid directly to the dewar, and regulates to the amount of helium in the dewar. The quantity of helium is often measured by the liquid level probe, LL11, which expresses itself in units of "% full." Sometimes it is measured by another type of probe, DP11, which measures the level in the less intuitive unit of "inches of water;" 4.6 inches represents a dewar which is 100% full. The dewar should be about 60% full, so EVLH is opened if the liquid level drops below that level.

EVQH (Loop 15) adds liquid helium to the exhaust of the wet engine (or JT valve) before it is heat exchanged with the helium in the dewar. As such, it does not directly contribute to the dewar level. But indirectly it does, because it means that the 2-phase helium will be colder (see below). EVQH is normally used only during quench recovery. In the original design, the feedback loop for EVQH looked at TI7, which measures temperature at one point inside the heat exchanger. In practice, it now often uses the liquid level of the dewar, with the set point at 50% full (10% lower than the set point for EVLH, so they don't fight each other.)

2-Phase Helium

After the 2-phase helium has been used for heat exchange in the magnets and feed can, it enters the dewar. The liquid helium settles at the bottom of the dewar and the cold gas moves on to the cold compressor and the heat exchanger. It is important that no liquid helium gets inside the cold compressor, as explained in the next section.

Cold Compressors

During "normal" operations, the cold helium liquid in the dewar evaporates and becomes the cold helium gas used as counterflow in the heat exchanger. During low-temperature operations, the helium from the dewar first passes through the *cold compressor,* and then continues on through the heat exchanger. The cold compressors should not be confused in any way with the ring compressors!

Since the helium at this point is already quite cold, and since it has already been through the magnets, it may seem counterintuitive to put a compressor into the system knowing that it will heat the gas about to enter the heat exchanger. Indeed, nearly half of the heat load of the system comes from this compressor, and compensating for that heat presents a new engineering challenge. The advantage actually comes from lowering the pressure of the *input* of the compressor, that is, the 2-phase helium. Since the 2-phase circuit is continuous all the way back to the magnet JT valves, lowering the pressure lowers the temperature throughout the system. The drop of about 0.8K is enough to allow the magnets to sustain currents and fields equivalent to 1 TeV.

The compressors are of a turbine design, with the blades spinning at a minimum of 20,000 RPM (20 KRPM) and a maximum of 95 KRPM. The faster they spin, the lower the input pressure, and the greater the heat load on the system. (At those speeds, a tiny droplet of liquid could create enough centrifugal force to destroy the compressor.) Although most of the liquid of the 2-phase condenses in the dewar, there are heaters at the input of the compressor to take care of any leftover mist.

The compressor speed regulates to PI11 (Loop 16), a pressure gauge located at the point where helium is taken from the compressor. The set point, when the compressors are fully operational, may be around 11 psia. Notice that the units are in *absolute* pressure and that the operating point is sub-atmospheric. All of the 2-phase circuits between the magnet JT's and the cold compressor had to be "hardened," to prevent leaks, before adding the cold compressors to the system.

The lower pressure in the 2-phase circuit creates a larger pressure differential with respect to the single phase, which increases the flow across the magnet JT valves. The valves are closed down by a few percent in order to reduce the flow to normal levels. The increased "JT effect" also helps to lower the temperature.

There is a valve, PVCB, which bypasses gas around the compressor. (The "P" means that the valve is pneumatically driven.) If left in remote, the valve closes automatically when the compressor turns on.

It should now be possible to explain why the dewar is a necessary adjunct to the cold compressor. The exhaust from the compressor is warmer than the gas that would otherwise be returning through the heat exchanger. Incoming warm gas from the exchanger works its way through the wet engine, producing a lower quality of cooling than would otherwise be the case. The dewar is an opportunity for the wet engine exhaust to heat exchange with a colder substance before it enters the magnets.

The cold compressors are not turned on until the refrigerator has already been cooled to "normal" cryogenic temperatures.

*　　*　　*　　*　　*

Finally, the exhaust, whether or not it has been through the cold compressor, passes through the heat exchanger and absorbs all of the heat it can from the incoming high pressure helium. (In an ironic twist of fate, a few hours earlier it *was* the incoming helium.) Now, at roughly room temperature, it enters the ringwide suction header and is taken in by the ring compressors, beginning the cycle again.

Power Leads

Another cryogenic component is required before the Tevatron is allowed to ramp—*power lead* cooling. The power leads are where the copper bus from the power supplies meets the superconducting alloy in the feed can. The copper, of course, is not superconducting, but it is required to carry several thousand amps to the feed can—there is considerable resistive heating at the leads. Leads at the "2" and "3" locations are associated with the feed cans, because there are power supplies at those buildings. At the "1" buildings they are at the upstream end, because they link to the copper bus of the straight-section bypass; at the "4" buildings they are at the downstream end for the same reason. Helium is tapped off of the single-phase circuit and passes through the superconducting part of the leads in order to cool them; the copper is cooled with LCW.

Current in the Tevatron, and therefore the cooling requirement, varies with operational demand. Sometimes the Tevatron is off, sometimes it is ramping, sometimes it is in a 150 GeV store and sometimes a 1 TeV store. Cooling which would be appropriate at 1 TeV would freeze the water going through the leads during a 150 GeV store.

Control of the flow is implemented via three orifices upstairs in the refrigerator building, two of which are equipped with solenoids. The two with solenoid control can be interchanged, if necessary, with orifices that have different diameters and flow rates; the remaining orifice is fixed.

Helium flows through the fixed orifice at 15 SCFH (Standard Cubic Feet per Hour). It is always open.

The second orifice, solenoid-controlled, carries either 40 or 60 SCFH of helium. Since the solenoid is open whenever the Tevatron is above 70 GeV or so, it is known as the *low-energy* solenoid.

The third orifice, with a flow of either 60 or 80 SCFH, has a solenoid which is opened whenever the TeV current is above 150 GeV. It is known as the *high-energy* solenoid. Actually, all three branches have to be open when the Tevatron is at higher energies.

A platinum resistor measures the temperature of the helium downstream of each lead. (There are "upstream" and "downstream" leads, although the terminology is strictly true only for the "2" and "3" houses; both power leads for the "1" houses are at the upstream end and both leads for the "4" houses are at the downstream end.) The parameter for the upstream lead temperature is TPUL, and that of the downstream lead is TPDL.

Using a slightly confusing convention, digital control for the solenoids uses TPUL and TPDL for parameter names, but assigns the low-energy solenoid to TPDL and the high-energy solenoid to TPUL. In reality, "low-energy" and "high-energy" have nothing to do with "upstream" and "downstream" when it comes to digital control of the leads.

Control of the lead flows is through parameter names (they are listed on page F14, subpages 33 and 34), the FSM, or the Sequencer.

The FSM automatically turns on the high-energy flows whenever the Tevatron energy exceeds 160 GeV (i.e., the Tevatron starts ramping). (It is preferred, however, not to risk a ramp dump by forgetting to turn them on manually.) They will also turn off the high-energy flows if the Tevatron has not been above 160 GeV for 5 minutes, in order to prevent the leads from freezing. (There are fans downstairs blowing on the leads to minimize the risk of freezing, but again, it is preferable not to take chances.)

Finally, the temperature readbacks for TPUL and TPDL are monitored by the FSM to make sure that the leads are not getting too hot. The limit is normally 297K. If the limit is exceeded, the FSM removes the ramp permit. (The ramp permit looks at several different measurements and is described in a section below.)

Safety Leads

At every other spool piece, *safety leads* are attached to the bus. The safety leads connect the bus to the QBS switches, and carry the current away from a quenched cell. If a particular cell quenches repeatedly, the leads could overheat. Cold helium is pulled from the magnets to prevent that from happening. Safety lead temperatures are not monitored from the Control Room, but there are flow meters in the tunnel which must be checked if the quenches become too frequent.

The Cooldown Sequence

This section will describe the cooldown sequence beginning at liquid nitrogen temperature, which is about 77K. Cooling from room temperature to nitrogen temperature is the task of specialists for whom the information in this chapter would be superfluous.

If the magnet string is sufficiently warm, four *modes* of cooling are required—*cooldown, transition, fill,* and *operate.* The control of the sequence can be done manually through the feedback loops, or by the Finite State Machine (FSM). Understanding the feedback loops is an important prerequisite for understanding what the FSM does; the emphasis will be on the loops for this discussion.

Cooldown

This stage (which has admittedly been given a totally ambiguous name—isn't it all cooldown?) refers to the first surge of cold helium passing through the magnet string. The magnets are too warm at this point to allow any of the helium to remain as a liquid; also, the single-phase circuit itself is likely to be full of warm helium. Cold helium enters from the feed can and makes one pass through the single phase circuit, displacing the warm gas. The helium being displaced is too warm for the magnet JT valve to be of any use; in fact, above 35K or so, JT expansion actually *heats* the gas. Instead, displaced helium is dumped directly into the suction header via the *cooldown* valves at the end of the string. If the refrigerator is strong enough to sustain the load, Kautzky valves can be used to speed up the process.

At first glance, the feedback loops for the cooldown valves (Loops 13 and 14) seem rather counterintuitive (i.e., they don't make sense). They look at TR12, a carbon resistor at the inlet of the wet engine. Because of its location, TR12 is a good indicator of heat exchanger performance.

What is keeping the heat exchanger cold? Remember that during normal operations, the counterflow of cold gas comes from returning 2-phase helium. However, if the magnet

JT valves are closed and the cooldown valves are open, there will be no returning cold helium. This problem is solved by changing the set point of the wet engine to 20 psig, while leaving the set point of the bypass valve at 18 psig. The wet engine will speed up to its maximum of 1200 RPM as it tries to raise the pressure at PI13 to 20 psig, but the bypass valve will open at 18 psig and shunt the extra helium into the dewar. That cold helium is what finds its way into the heat exchanger.

If the cooldown valves are open too far, the pressure will drop throughout the single phase circuit, and the bypass valve won't shunt enough cold helium into the dewar. The heat exchanger, and therefore TR12, will consequently warm up. To preserve the quality of the cooldown wave, the feedback loops will close down on the cooldown valves to raise the pressure in the circuit.

The set point of the loop during this stage is usually set at either 10K or 20K, and produces a "10˚ wave" or a "20˚ wave." The 10˚ wave is slower, because the valves are not as far open, but it produces a better quality of cooling for this stage.

Often, the Kautzky valves are opened during cooldown mode. Any Kautzky downstream of the wave can be used. Since the Kautzkys were originally designed as relief valves, the volume of helium purged from the string is enormous compared to the cooldown valves. They are only opened for a few seconds at a time; TR12 rises very rapidly during that time. They are then closed until TR12 indicates that it is safe to proceed. There is no loop control for the Kautzky valves; they are controlled either manually or through the Finite State Machine.

EVX1, which allows incoming helium to be exchanged with liquid nitrogen, is opened during cooldown mode to assist the heat exchanger. (It is normally closed to a trickle during steady-state operations.)

EVQH, usually closed during normal operation, is opened, and regulates to the liquid level of the dewar. During cooldown phase the dewar is not likely to fill, so EVQH remains open at its maximum value. Officially, EVLH (normally open) is closed, although in practice it is sometimes opened if needed. (EVLH would try to fill the dewar directly, but it is not necessary for the dewar to have liquid during this phase. It can, however, keep the heat exchanger colder.)

Transition

The cooldown phase is terminated when the cooldown wave reaches the end of the string. At that time, the magnet JT valves are opened to 100% and the cooldown valves are closed. It is important to properly assess when and where that happens—if the change is made too soon, helium going through the JT valve will be too warm and the counterflow will actually heat the single phase instead of cooling it. If it is done too late, the cold helium will be dumped into the suction header and wasted; also, since the cold gas has a higher density, flow through the system will increase, helium will leave faster, and less will be available to cool the heat exchanger.

The switch to transition mode should be made when the end of the string reaches 9K or so; if the temperature is being plotted, the change can be delayed until the drop in temperature begins to level off. The cooldown and magnet JT valves are located in the turnaround box at the end of the string, so it is important to look at a carbon resistor in or near the turnaround box (Fig. 5.2). The upstream measurement at "2," "3," and "4" buildings is read from TRQ1. The "1" houses must use TR15 for the upstream measurement because the lattice requires additional quadrupoles upstream of TRQ1. At B1 and D1 it is the low beta quads, and at A1, C1, E1, and F1 it is the 82" and 99" lattice-matching quads that separate TR15 and TRQ1.

Downstream, the "2," "3," and "4" houses use TRQ9 as the end of the string, except for A4 and C4, which use TR18 because of the low beta quads. The downstream "1" strings extend to the "21" location, so they use TRQ0.

The reason that TR15 and TR18 are not always used is that they can be susceptible to "cross-talk" from adjacent turnaround boxes. For example, if the downstream F2 string is warm, F3TR15 may read artificially high.

The magnet JT is locked open at 100% during transition because at the temperatures typical of this stage the JT effect is relatively weak, and *flow* is favored over restriction of the valve. The flow/restriction ratio will change as cooling progresses, but at this point establishing a reasonably cold counterflow in the 2-phase circuit, and lots of it, is the most important goal.

Fill Mode

To "fill" the magnets is to coax the helium to condense into a liquid. The change from transition to fill mode is implemented when the end of the magnet string reaches 5.3K; the magnet JT is changed from 100% open to 80% open, since the helium is now cold enough to take greater advantage of the JT effect. Also, the set point for the wet engine is returned to its operational value of 17 psig; the counterflow through the 2-phase circuit is sufficiently cold to allow the bypass valve to be closed. The reduced pressure in the circuits will also aid in the cooling, although the engine will continue to run fairly fast as long as the JT is open at 80%.

Initially, the DT will be high—around 6 or 8 psid. After what may seem like an eternity, liquid will begin to condense in the 2-phase circuit and the DT will begin to come down. The dewar will begin to collect some liquid helium as well. The magnet JT valves should be closed down very gradually during the latter stages of fill mode.

Operate mode

At operating temperatures, the helium is cold enough to take full advantage of the JT effect. EVLH is opened to maintain liquid level in the dewar, and EVQH is closed. EVX1 is reduced to a trickle, nitrogen no longer being needed by the heat exchanger. All of the loops can be restored to their operational steady-state values.

Low Temperature Operations

If high-energy beam is desired—which should be the case throughout the Collider run—the cold compressors must be used. The cold compressor is left off during most of the cooldown, because of the strain it puts on the heat exchanger. Only after a "normal" permit is established is the cold compressor turned on.

The purpose of the cold compressor is to lower the 2-phase pressure, as measured by PI11. The faster the compressor spins, the lower PI11 will be. The lower pressure makes the magnets colder, but there are two negative consequences: (1) the return gas flowing through the heat exchanger is warmer, and (2) there is greater flow through the magnet JT's, because of the increased differential pressure across the valves. For those reasons, the cold compressor is initially turned on to a speed of 20 KRPM; the set point of PI11 starts out relatively high. The cold compressor speeds up as the set point of PI11 is gradually lowered; at the same time, the JT valves are slowly closed down in order to reduce the flow. The 2-phase gradually becomes colder as the heat load from the compressor increases.

The sequence of cooling from "normal" permit temperatures to high-energy temperatures is coordinated through a Finite State Machine; it requires about half an hour to execute.

Refrigerator Ramp Permits

One of the tasks of the local controls system—specifically, the FSM—is to determine whether or not the system is healthy enough to allow the Tevatron to ramp. If it is, it issues a permit to the local QPM. The QPM's report to TECAR; there must be permits from all 24 refrigerator buildings before the Tevatron is allowed to ramp.

The permit is implemented via a ramp *inhibit*, not a ramp dump. For example, if a refrigerator parameter drifts out of tolerance at flattop, the Tevatron finishes its current ramp normally but locks at 90 GeV when it reaches that level. In Collider mode the refrigerator permit is completely ignored once a store is established. Given the investment required to set up collisions, it is better to risk a quench than to guarantee the death of a store.

There are two types of permits—a *normal* permit and a *low-temperature* permit. The normal ramp permit is required for any current up to 900 GeV. The low-temperatutre ramp permit, which includes everything required of the normal permit (and more), is necessary for operations above 900 GeV.

The ramp permits should not be confused with the refrigerator alarms on AEOLUS, which may or may not have the same limits for a given device.

The normal permit looks at parameters involving pressure, temperature, and DT's.

- *Pressure:* PI13, PI14, and PI17 (helium pressure upstream of the feed can, upstream string, and downstream string, respectively) are monitored to determine if they are within minimum and maximum limits. The FSM customizes the limits for each device, but a typical minimum value is 14 psig (remember that during normal operations, the wet engine set point is 17 psig). Common reasons for low pressure might be (1) the magnet JT valves are not close enough to their minimum value following a cooldown, (2) there is a leaky Kautzky valve, or (3) there is a problem with the wet engine or pressure transducer.

- *Magnet temperature*: Most of the carbon resistors, including those in the feed can and all of the TRQ's, are monitored for temperature violations. A typical limit is 5K, but resistors in the feed can and turnaround boxes are allowed more leeway. If the cooldown sequence is complete and the temperatures are close to the limit, and the permit is still missing, opening the magnet JT slightly or turning on the low-energy lead flows might bring it back.

- *Lead flow temperature:* The limit for the lead flow temperatures is normally 297K before the ramp permit is removed. (This *temperature* permit, like the other permits listed here, is a subset of the refrigerator permit sent to the QPM—it should not be confused with the *voltage* monitoring of power leads done by the QPM itself. The penalty for a voltage violation is a ramp *trip*, not the ramp inhibit of the refrigerator permit.)

- *DT's:* Violation of a DT limit does not immediately revoke the ramp permit, because it usually takes a few minutes for the lack of subcooling to impact the temperature inside the magnets. Instead, when a DT rises above its set point, a "counter" is initiated; if the DT is slightly above the set point the numbers accumulate slowly, but if the DT is significantly high they accumulate quickly. When the integrated value reaches a value 1300, the permit is removed. The number is reset to zero as soon as the DT drops below the set point. One technique for deferring a ramp inhibit—to be applied judiciously—is to momentarily put the set point at an unrealistically high value in order to reset the counter.

- The *low-temperature permit* includes all of the requirements of the normal permit, and adds one. PI11, which measures the inlet pressure of the cold compressor, must be below the set value. PI11 is an indicator of pressure throughout the 2-phase circuit, and the magnet temperatures rise and fall with that pressure. The value of the set point is chosen to produce the desired temperature.

Refrigerator Controls

Control of the Tevatron cryogenic systems is implemented through VME crates distributed around the ring. There is one VME crate per ring sector, located at the zero buildings.

Each VME uses an ARCNET loop to control its little empire. (A strand of the 19-conductor cable was cut into six sections to create the loops.)

Locally, at the buildings, the commands and readbacks (i.e., I/O) are interfaced to the hardware through Intel 186 processors. There are actually two systems at the "1" through "4" houses: the *device* I/O crate and the *thermometry* I/O crate; each system has its own 186 card. The compressors are controlled by a device crate, but they do not need a thermometry crate because there are no low temperatures there to measure.

The device I/O crates are located in the refrigerator buildings; the crate interfaces with the valve actuators, engine controllers, and the cold compressor. There are also digital input cards and A/D converters for monitoring the instrumentation (e.g. the pressure transducers), and a vacuum card. (The vacuum card monitors conditions in the equipment upstairs, such as the valve box and U-tubes. It should not be confused with the CIA crate in the service building, which spies on the vacuum in the beam pipe and magnet cryostats.)

The thermometry crates are responsible for providing the carbon and platinum resistors with the pulses of current necessary to measure the resistance, and therefore temperature. They are located in the Tevatron service buildings, on the back wall behind the QPM racks. In addition to their role of measuring temperature, the thermometry crates have a link to the QPM's.

All of the software, whether for thermometry or the device crate, gathers data at a 1 Hz rate. Altogether, there are about 700 readbacks from each house.

Loop Control

The reader who has successfully followed the preceding discussion of refrigerator operation already has a good intuitive understanding of how the loop controls work. A few points should be made anyway about F8, the loop control page. There are 16 loops normally used in refrigerator control. Remember that the loops and the Finite State Machines are two distinct entities; failure to recognize this difference is a common source of confusion for newcomers.

Itemizing some of the parameters on page F8:

- *Input variable,* or the parameter that the loop is trying to maintain. Typical input variables include temperature or pressure readbacks. The purpose of the loop is to maintain the parameter at the *set value.*
- *Output variable,* or the device responsible for maintaining the parameter at the set value. The device associated with each loop number is constant for all operational Tevatron refrigerator buildings and might as well be listed: (1) EVBY, (2) SPWE, (3) EVUH, (4) EVDH, (5) EVJT, (6) SPWE, (7) EVX1, (8) EVX2, (9) EVLN, (10) EVLH, (11) EVUN, (12) EVDN, (13) EVUC, (14) EVDC, (15) EVQH, and (16) SPCC. Details on each of these is scattered throughout the text above.
- *Minimum and Maximum Positions:* ("Position" can also mean things other than valve position, such as engine speed.) When the loop is enabled and active, the device is

free to roam within these limits in order to bring the input variable to its set value. If a position outside of the range is needed, the loop is helpless.

- *Current Position:* A readback, but it is also possible to type in a desired position.
- *Enable/Disable:* When disabled, the position is locked in place and will not respond to loop controls. However, it is still possible to type in a desired position, and the Finite State Machine can still change the position.
- *Active/Inactive:* This is a tricky one—it indirectly reflects what the Finite State Machine is doing. When the loop is "Inactive," the FSM is actually controlling the device, based on algorithms that have nothing to do with what is on the screen. Moreover, one can toggle the bit to "Active," but the FSM continues its work in secret.
- *Remote/Local:* This is controlled by the switch on the actuator card in the Device I/O crate.
- *Sample Time:* A measure of how often the loop checks conditions, and therefore how quickly the device changes. There is a wide range of response times. For example, EVBY must respond rapidly to pressure fluctuations, and has a sample time of 6 seconds. EVUN and EVDN, which control flow through the nitrogen shield, respond to slow temperature fluctuations and have a sample time of over 15 minutes.

Finite State Machines

The *Finite State Machines (FSM's)* provide automatic control for various cryogenic systems, including quench recovery and cooldowns. Each house can run up to 32 FSM's; collectively, they can be thought of as a controls system in parallel with the loops. When called upon, the FSM's will inactivate selected loops and take over their functions until the task is completed. (The FSM's can be very efficient, as long as cryogenic conditions are predictable, but—just in case—human oversight should never be neglected.)

The FSM's can perform the following tasks:

- Quench response. Coordinating the quench response is the responsibility of the QPM, but the QPM makes a request that the FSM begin refrigerator recovery. A QPM failure will also trigger a frig recovery, unless the response is disabled (via page F3).
- Automatic cooldowns, which can also be initiated from page F3. A separate FSM is needed to turn on the cold compressors for low temperature operations.
- Control of lead flows; the FSM continuously monitors the lead flows, turning them on when the Tevatron begins to ramp, and turning the appropriate ones off if the TeV is off or set at 150 GeV.
- Monitoring conditions for a ramp permit. For example, if a temperature wanders out of limits, the FSM tells the QPM that the permit is being revoked, and the QPM will tell TECAR to hold the ramp at 90 GeV.

In the case of quench recovery and automatic cooldowns, the FSM's begin by evaluating current conditions—such as temperature, pressure, or liquid level in the dewar—and then executing a predetermined set of algorithms to achieve the desired result. At each stage, called a *state*, the FSM monitors a specific set of conditions. When those conditions are met, the FSM advances to the next state (fortunately, there are only a finite number of states). Each state has a "library" of *operations, actions,* and *timers* it calls upon to do its job and evaluate its own performance. The FSM's, states, operations, actions, and timers are all explicitly listed on page F13, although it can be difficult to sort them all out from that page. For a more intuitive interpretation, live graphic displays for cooldowns can be launched from page F24, and displays for the refrigerator permits can be launched from F23.

The Consolidator

The "Fridge Consolidator," as it is affectionately known, is housed in a VME crate in the Computer Room. Its purpose is to gather data from around the ring and organize it for display applications. Programs dependent on the Consolidator include F3, F5, F16, and F17, along with any other application which aspires to present a ringwide view of the refrigerators. Failure of the Consolidator results in no data being returned through these programs.